

This Page Is Inserted by IFW Operations
and is not a part of the Official Record

BEST AVAILABLE IMAGES


Defective images within this document are accurate representations of the original documents submitted by the applicant.

Defects in the images may include (but are not limited to):


- BLACK BORDERS
- TEXT CUT OFF AT TOP, BOTTOM OR SIDES
- FADED TEXT
- ILLEGIBLE TEXT
- SKEWED/SLANTED IMAGES
- COLORED PHOTOS
- BLACK OR VERY BLACK AND WHITE DARK PHOTOS
- GRAY SCALE DOCUMENTS

IMAGES ARE BEST AVAILABLE COPY.

**As rescanning documents *will not* correct images,
please do not report the images to the
Image Problems Mailbox.**




- (72) NAQVI, Shamin Abbas, US
(72) IMIELINSKI, Tomasz, US
(71) BELL COMMUNICATIONS RESEARCH, INC., US
(51) Int.Cl.⁶ G06F 17/60
(30) 1995/12/08 (08/569,639) US
(54) **PROCEDE ET DISPOSITIF DE DIFFUSION DE MESSAGES
PUBLICITAIRES SUR UN RESEAU INFORMATIQUE**
(54) **METHOD AND SYSTEM FOR PLACING ADVERTISEMENTS IN
A COMPUTER NETWORK**



(57) Avec ce procédé et ce dispositif, un serveur disposant d'un ensemble de messages publicitaires est mis en communication électronique avec un réseau informatique. Les messages de ce serveur ne sont pas liés à une page particulière contenant des informations disponibles sur ce réseau, mais ils arrivent en réponse à une requête introduite par l'utilisateur (17) puis sont mélangés de façon dynamique avec le contenu des pages

(57) A method and system for placing advertisements in a computer network. A server containing a collection of advertisements is placed in electronic communication with the computer network. The advertisements on the server are not tied to any particular page containing information on the network, but rather, are retrieved in response to a query entered by the user (17) and dynamically mixed with the content of the pages



(21) (A1) **2,240,009**
(86) 1996/12/06
(87) 1997/06/12

envoyées en réponse à cette requête (16). On peut ainsi afficher des pages de contenu assorties de messages publicitaires ciblés et spécifiques intégrés à une page, selon une disposition particulière. Ces messages peuvent être adaptés à une série de contraintes (19) demandées par l'utilisateur ainsi qu'à celles de l'éditeur de cette page. Ce dispositif fait appel à des contrats (21) précisant les règles de commercialisation (18) qui lient des messages publicitaires à des requêtes spécifiques, ce qui permet aux publicitaires de cibler un public précis et de garantir un certain degré de diffusion aux messages dans l'espace publicitaire le plus favorable. Des algorithmes permettent de contrôler la cohérence contractuelle, ce qui permet de veiller à ce que tous les contrats acceptés soient correctement appliqués. On décrit aussi une fonction d'affichage publicitaire pendant l'attente d'un téléchargement.


returned in response to the query (16). The present invention displays the content pages with focused, targeted advertisements as a part of the page, in accordance with a particular layout. The advertisements can be made to satisfy a set of constraints (19) requested by the advertiser, as well as the constraints of the publisher of the page. The system uses contracts (21) to specify the marketing rules (18) that link ads with specific queries, to permit advertisers to target a specific audience, and to guarantee a certain amount of exposure of the advertisement in prime advertising space. Algorithms are used to check for contract consistency to ensure that all contracts that are accepted can be properly satisfied. The present invention also provides a download delay-time advertising feature.

**PCT**WORLD INTELLECTUAL PROPERTY ORGANIZATION
International Bureau

INTERNATIONAL APPLICATION PUBLISHED UNDER THE PATENT COOPERATION TREATY (PCT)

(51) International Patent Classification ⁶ : G06F 151/00, 153/00		A1	(11) International Publication Number: WO 97/21183
			(43) International Publication Date: 12 June 1997 (12.06.97)
(21) International Application Number: PCT/US96/19509 (22) International Filing Date: 6 December 1996 (06.12.96) (30) Priority Data: 08/569,639 8 December 1995 (08.12.95) US (71) Applicant: BELL COMMUNICATIONS RESEARCH, INC. [US/US]; 445 South Street, Morristown, NJ 07960-6438 (US). (72) Inventors: NAQVI, Shamin, Abbas; 19 Spring Valley Road, Morristown, NJ 07960 (US). IMIELINSKI, Tomasz; 153 Willowbrook Drive, North Brunswick, NJ 08902 (US). (74) Agents: GIORDANO, Joseph et al.; International Coordinator, Rm. 1G112R, 445 South Street, Morristown, NJ 07960-6438 (US).			(81) Designated States: CA, JP, KR, European patent (AT, BE, CH, DE, DK, ES, FI, FR, GB, GR, IE, IT, LU, MC, NL, PT, SE). Published With international search report.

(54) Title: METHOD AND SYSTEM FOR PLACING ADVERTISEMENTS IN A COMPUTER NETWORK



(57) Abstract

A method and system for placing advertisements in a computer network. A server containing a collection of advertisements is placed in electronic communication with the computer network. The advertisements on the server are not tied to any particular page containing information on the network, but rather, are retrieved in response to a query entered by the user (17) and dynamically mixed with the content of the pages returned in response to the query (16). The present invention displays the content pages with focused, targeted advertisements as a part of the page, in accordance with a particular layout. The advertisements can be made to satisfy a set of constraints (19) requested by the advertiser, as well as the constraints of the publisher of the page. The system uses contracts (21) to specify the marketing rules (18) that link ads with specific queries, to permit advertisers to target a specific audience, and to guarantee a certain amount of exposure of the advertisement in prime advertising space. Algorithms are used to check for contract consistency to ensure that all contracts that are accepted can be properly satisfied. The present invention also provides a download delay-time advertising feature.

**METHOD AND SYSTEM FOR PLACING
ADVERTISEMENTS IN A COMPUTER NETWORK**

NOTICE OF COPYRIGHTED MATERIAL IN DISCLOSURE

A portion of the disclosure of this patent document contains material which is subject to copyright protection.

5 The copyright owner has no objection to the facsimile reproduction by anyone of the patent document or the patent disclosure, as it appears in the Patent and Trademark Office patent file or records, but otherwise reserves all copyright rights whatsoever.

10 **BACKGROUND OF THE INVENTION**

Field of the Invention

The present invention relates generally to advertising and, in particular, to a method and system for placing advertisements in a computer network, such as a wide area
15 network or a metropolitan area network.

Description of the Prior Art

A number of information services are currently offered on wide area public networks, such as the Internet. The
20 basic model of these services is that a query is processed and a set of answers is returned. For example, a number of yellow page searchable directories have been developed in which users can query categories of products or services (e.g., new car dealers) and the system returns a set of
25 vendors satisfying the query.

One problem with these services is determining who should subsidize the end user's cost of using the services.

One way of providing such a subsidy is through the use of

advertisements. Public networks, such as the Internet, are starting to see random use of advertisements. For example, some search engines used in public networks are sponsored by a single vendor or company. The sponsoring vendor or company will often have an advertisement displayed on a user's computer screen during use of the search engine. Other search engines display randomly selected advertisements from a fixed set of advertisements.

As one example, a number of services on the World Wide Web ("WWW") are currently using ad hoc means of advertising. That is, when a user uses certain search engines for conducting a search, the user will be shown advertisements while doing the searching. These advertisements are sometimes referred to as "banner" advertisements because they simulate a banner that the user sees as the user is traveling down a "road" on the computer network. These advertisements are typically tied to a particular search page that the user encounters during the search.

The current state of the art is such that when the user uses a search engine, a randomly selected advertisement is shown as if it is part of the search page.

For example, the user may enter a search request to see a home page on cooking and, as a part of that page, the existing systems might display an advertisement about cars.

This is a problem, of course, because there is no connection made between the content of the advertisements or the message of the advertisements and what the user is actually searching. Another problem with the current state of the art is that there is no guarantee associated with the advertisements that they will be shown a certain number of times or in a certain place on the page with

respect to the other pieces of data on that page. For these reasons, the existing systems for placing advertisements on computer networks, such as the Internet, fail to take full advantage of the interactive/demand driven environment of the online medium.

SUMMARY OF THE INVENTION

It is thus an object of the present invention to provide an improved method and system for advertising on a computer network, particularly a wide area or metropolitan area network, which overcomes the above described problems of existing advertising systems.

It is a further object of the present invention to provide an effective advertising medium that anyone with a personal computer can use to advertise on computer networks.

It is a further object of the present invention to provide a method and system for advertising in which advertisements are more informative and provide a direct link to additional information about the advertised product or service or about the advertiser.

It is a further object of the present invention to provide a method and system for advertising on a computer network in which advertisements are more focused and targeted, for example, by user queries and user profiles, including the past history of the user's interactions with the system.

It is a further object of the present invention to provide a method and system for advertising in which new online publishers can be created with little effort.

It is a further object of the present invention to provide a method and system for advertising in which

advertisers can be guaranteed that their advertisements will be displayed a certain number of times or in a particular manner or under particular circumstances.

5 Additional objects, advantages and novel features of the invention will be set forth in the description which follows, and will become apparent to those skilled in the art upon reading this description or practicing the invention. The objects and advantages of the invention may be realized and attained by the appended claims.

10 The present invention provides a new process and system for online advertising. This new process will be referred to throughout this application as query-based advertising ("QBA"). In the QBA process, advertisements are primarily triggered by user queries. User queries, as
15 used herein, refer to requests from an information consumer for one or more pages of information from a computer network. As a result of a query, a user is exposed to advertisements with the present invention, i.e., the query triggers advertisements.

20 In accordance with the present invention, a server containing a collection of advertisements is placed in electronic communication with a computer network. The computer network is preferably a wide area network, such as the Internet, or a metropolitan area network. As used in
25 this application, the phrase "computer network" refers to any public or private data communications network.

The advertisements on the server are not tied to any particular page containing information on the computer network. Rather, the advertisements are contained on the
30 server, distinct from the pages that may or may not later carry the advertisements. The pages by themselves have no

advertisements. Thus, the pages are analogous to a newspaper or magazine devoid of any advertisements.

When the user requests a certain page or a certain topic of information, the relevant pages are retrieved from the computer network and shown to the user. The present invention, upon receiving the user's request, retrieves advertisements that are related to the user's action, dynamically mixes the advertisements with the content of the pages according to a particular layout, and displays the pages with focused, targeted advertisements as a part of the page. The advertisements can be made to satisfy a set of constraints requested by the advertiser, as well as the constraints of the publisher of the page, as further discussed below.

The advertisement triggering mechanism of the present invention is not random or coincidental, but rather, is prespecified in advance. This specification will be referred to in this application as a contract. A contract specifies the marketing rules that link advertisements with specific queries. For example, a diet soft drink advertisement may be shown when a user asks for a page about exercising equipment. These rules are specified by advertisers implementing the concept of "focus" or

"relevance" of advertisements and help the advertisers to target a specific audience. Owners of pages specify the focus content of their pages through special tags within a page. These tags are not displayed to the information consumer; the tags are used to decide what advertisement can be shown when the page is requested by a consumer.

The notion of a contract, however, goes well beyond just marketing rules. First of all, the advertising space on the online medium, although technically unlimited, is

severely restricted by the user's attention span. Placing advertisements on the first page which constitutes the answer to a query gives the advertisements much higher probability to be seen than on later pages of the answer.

5 This is analogous to a certain extent, to newspaper advertising where advertising on the first page of the newspaper is more expensive than on the last page of the newspaper. Thus, the present invention introduces the concept of prime space as a part of the QBA.

10 Prime space will be used in this application to refer to limited advertising space on the first page of the answer to a query. Prime space is, in principle, as attractive as prime time television. Prime space, like prime time, is limited so it may be impossible to show all
15 the relevant advertisements in prime space.

As a solution to the prime space limitations, the present invention provides guarantees to advertisers through the use of contracts. The contracts provide tools to specify such guarantees and to enforce them.

20 Specifically, the present invention provides the following types of guarantees (or "contracts") to advertisers:

1. Ratio-based: The ad will be shown in prime space no less than a certain percentage of the times a relevant query is made by a searcher.
- 25 2. Exclusive or Competitive: Ad A is never to be shown in prime space when ad B is shown.
3. Dependent: Ad A is to be shown with a predetermined probability when ad B is shown in prime space.
- 30 4. Story: An ad consisting of several sub-ads will be shown in prime space in sequence to a consumer.

5. Frequency-based: The ad will be shown in prime space a certain number of times in a given time period (e.g., 100 times per month).

5 Additionally, the contracts may also specify the type of layout used to mix ads together with the data. This may include the layout types, and so forth.

10 The present invention also provides means for determining whether there are any contracts that cannot be satisfied. For example, if the prime space is sufficient to hold only two ads, it is impossible to guarantee 100% ratio for more than two ads. Similarly, if ad A is never to be shown when ad B is shown and if both ad A and ad B are equally relevant according to focus matching, then it is impossible to guarantee ratios above 50% for both ads A and B.

15 The present invention uses algorithms to check for contract consistency to ensure that all contracts that are accepted can be properly satisfied. An advertiser can ask for a certain contract (e.g., 5% relative coverage in prime space) and the present invention will determine whether that contract can be satisfied. If not, the system suggests what coverage it can provide. After the contract is determined to be consistent, the present invention will enforce the contract by displaying advertisements in accordance with the contract.

20 A consequence of QBA is that ads cannot be placed on pages a priori because it is the query that determines what ads are to be placed on a page. This is referred to as dynamic advertising. The query asks for a page that has a focus. Ads that are resident in the system are checked to determine which ads can potentially be placed on the page in question. This decision is based on matching the focus

of a page with the focus of the ad. When not all matching ads can be placed on a page because of space limitations, the contract enforcement feature of the present invention ensures that the ads that are placed on the page are
5 consistent with the contracts signed by the system with the advertiser.

Another consequence of QBA is that it must be determined where to place an advertisement on a page. Currently, pages on the WWW have a fixed a priori agreed
10 upon area on a page for advertising. The present invention, however, dynamically determines where an advertisement should be placed. The present invention permits owners of pages to be free to organize the layouts of their pages. The page owner makes the final
15 determination of where an advertisement should or should not be placed. Once the page is organized, the page owner can specify where an advertisement can be placed by placing a special tag ("ad space") on their page. Like the focus tag, this ad space tag is not visible to the consumer. It
20 is a directive to the advertising system that the tag denotes space available for advertising, i.e., the tag can be thought of as a billboard stating "this space available for advertising."

The present invention also provides a download delay-
25 time advertising feature. Most consumers on computer networks, such as the Internet, experience a delay from the moment they ask for a page to the time the page is downloaded to their browser. This delay may be caused, for example, by traffic congestion on the WWW. This delay time
30 will be referred to as download delay-time.

The present invention uses the download delay-time to display ads to the consumer. This will be referred to as

down-load delay time advertising. The time that the consumer would have waited for a page to be downloaded is utilized to expose the consumer to ads. Once a page (or a subpart of the page) has been downloaded to the consumer's browser and is available to be displayed, the ads are replaced by the page content. The consumer does not experience a delay due to advertising by the present invention. Only the preexisting delay time is utilized by the advertising system.

In summary, the present invention provides a system and method for advertising on a computer network, comprising a server containing a plurality of advertisements, means for electronically connecting the server to a computer network, and means for selecting and retrieving an advertisement from the server in response to a query entered on the network. The selecting means comprises means for ensuring that a selected advertisement is relevant to the query. A mixer means is provided for combining a retrieved advertisement with a content page returned by the computer network in response to the query.

The mixer means comprises a layout manager means for computing an optimum layout of a combined page containing the retrieved advertisement and the content page. The mixer means also comprises a typography manager means for detecting special tags and HTML rules in the content page and for determining which part of the content page the selected advertisement can be displayed on. The content page is provided by a home page dispatcher, a search engine, or a generic HTML content provider in response to the query.

A transaction means is provided for permitting a user to make a transaction with the advertiser by interacting

with the displayed advertisement. A prime space manager means is provided for controlling when the advertisements contained on the server are selected by the selecting means, the control being based upon respective advertiser contracts associated with each of the advertisements.

An ad placement means is provided for placing new advertisements on the server, the ad placement means having means for entering an advertiser contract for each advertisement placed on the server and means for checking a new advertiser contract for consistency with existing advertiser contracts. A local client means is provided for displaying a retrieved advertisement only during a download delay time period in which a user is waiting for a content page to be downloaded in response to the query. A means for refreshing a displayed advertisement after a predetermined time period of inactivity by the user, and a means for replacing a displayed advertisement with a downloaded content page as soon as the content page is received by the client are also provided. The system includes means for calculating a search time estimation for the query, and means for combining a page containing the search time estimation and a retrieved advertisement. The system also includes means for determining a geographic location of a user, and means for selecting an advertisement based upon the user's geographic location.

BRIEF DESCRIPTION OF THE DRAWINGS

The present invention will become more clearly appreciated as the disclosure of the present invention is made with reference to the accompanying drawings, wherein:

Fig. 1 is a schematic diagram providing a general overview of the main components of the present invention.

Fig. 2 is a flow chart showing a process of placing advertisements on content pages provided by yellow page publishers of a computer network.

5 Figs. 3A-3D are screen prints showing page views displayed during various steps in the process of Fig. 2.

Fig. 4 is a flow chart showing the process performed by the layout manager of the present invention.

Fig. 5 is a flow chart showing the process performed by the parser of the present invention.

10 Fig. 6 is a flow chart showing the process performed by the mixer and ad selector of the present invention.

Fig. 7 is a flow chart showing a process of placing targeted advertisements in online pages provided by a home page publisher of a computer network.

15 Figs. 8A-8C are screen prints showing page views displayed during various steps in the process shown in Fig. 7.

Figs. 9A and 9B are flow charts showing a billing system process according to the present invention.

20 Fig. 10 is a flow chart showing the steps of an ad placement process according to the present invention.

Fig. 11 is a flow chart showing the steps of a page placement process according to the present invention.

25 Fig. 12 is a schematic showing the architecture of a transaction system according to the present invention.

Figs. 13A and 13B are screen prints showing page views displayed during various steps in a transaction process of the present invention.

30 Fig. 14 is a flow chart showing the steps of the transaction process according to the present invention.

Fig. 15 is a diagram chart showing the process performed by the prime space manager component of the present invention.

5 **DETAILED DESCRIPTION OF THE PREFERRED EMBODIMENT**

Reference will now be made in detail to a preferred embodiment of the invention, an example of which is illustrated in the accompanying drawings.

10 Fig. 1 provides a general overview or space architecture of the main components of the present invention. The layout manager 10 is primarily responsible for the computation of the layout of a page to be displayed to a client. Other modules of the system compute the necessary data that is to be displayed, as described below.
15 The layout manager 10 then accepts this data from the other modules and computes the optimum layout for the display.

20 The layout manager 10 uses two sub-components to determine the optimum layout; a rule base and a layout data base within the layout manager (collectively shown as the layout rules component 11 in Fig. 1). The layout data base contains a collection of templates that have been predefined and stored in the layout data base. The rule
25 base contains a collection of rules defined by the user or by the placer of the layout manager to ascertain which layout templates are best under what set of conditions. When the layout manager 10 receives the data to display, it invokes various rules in the layout rules to figure out which rule to trigger and, hence, decides on one layout as
30 being the best under the present set of conditions.

The variables used in the rules are the amount of data that is to be displayed, the kind of data that is to be

displayed, the amount of text, the amount of advertising material, and the size of the various maps and multi-media icons that are to be displayed. Based on these variables, the rules are checked and one or more rules are picked as delineating the best layout that can be found. The layout manager 10 then arbitrarily picks one of the rules that it thinks are the best and picks that layout, inserts the data given to it in the layout template, and gives the final product to the client to be displayed.

The variables that are expressed in the rules are based on a number of parameters including the size of the text, the amount of text, the size of the maps, the size of the images, and the size of the advertisement. In principle, there is no limit as to the number of parameters that one can achieve using the layout manager 10. The rules are preferably expressed in a programming language general enough to take any other parameters that one may desire to furnish.

The layout rules are constructed based on various parameters. The layout rules are then checked and used to pinpoint the exact template to be used for a certain kind of data situation. The template is then filled in by the layout manager 10 to provide a display to the client.

Next, the segmenter 12 will be described. The segmenter 12 is responsible for analyzing an HTML page and for extracting the relevant data from the HTML page. The purpose of the segmenter 12 is that, if the system is given a page constructed by another component or by some other content provider, the segmenter 12 will try to find reasonable places within that page where a certain advertisement could be inserted dynamically.

The segmenter 12 uses various rules for finding advertising space within a page, and for preventing the essential layout of the page from being disturbed. In particular, the segmenter 12 looks for certain special tags
5 inserted by the designer of the HTML page and replaces those special tags with advertisements or other types of information that it has been given. The segmenter 12 then looks at the page, decides where to insert ads, and provides that information to the layout manager 10.

10 In most instances, the typography manager 14 is the main input device to the layout manager 10. One exception to this will be described later. The typography manager 14's role is to "render" the page. By rendering the page it is meant that the typography manager 14 takes various
15 kinds of input that has to go on a page and decides which part of the information has to go on the page and which part of the information could in fact be deleted. Thus, the typography manager 14 is similar in function to a typographer in a traditional publishing situation.

20 The typography manager 14 has the final say in what data gets sent to the layout manager 10. The typography manager 14 also detects special tags and special HTML rules that tell it where the markups of the HTML page are supposed to happen. Ultimately, the layout manager 10 and
25 the typography manager 14 collaborate on coming up with a page. This page in its final form is then passed to the gate 15.

The gate 15 transfers a page from the system of the present invention to the client. The gate 15, therefore,
30 provides the gate between the client and the server. The gate 15 gives the page that it wants to display to the WWW

Daemon 16 and the Daemon interacts in turn with the client 17 on the user's side.

The above description has explained how the layout manager 10, the segmenter 12, the typography manager 14, and the gate 15 operate. Before describing in detail the remaining components of the system, an example flow of a request through the system will be explained.

Initially, a user requests a particular piece of information through one of the clients 17. The user's request is given to the WWW Daemon 16, which passes the information to the gate 15. The gate 15 at this point decides what piece of information is being requested by the user and finds other relevant pieces of information that can be commingled with what the user has asked. The user, for example, might ask the system to see certain car dealers, to find a phone number of a car dealer, or to get a page of a particular magazine.

The gate 15 at this point gives the request to the matching rule engine 18 ("MRE"). The purpose of the MRE 18 is to look at the content of the user's query and to find a category within its active index SIC 19 that matches the same type. If the user has asked for car dealers, the MRE 18 invokes its rules to determine that car dealers are part of a class of things relating to transportation. Based on the classification determined by the MRE 18, the system now knows that the user is asking about cars or about transportation or about whatever else that the user might be interested in.

The MRE 18 at this point then returns to the gate 15 the category index of the user's query. If the user had asked about cars or about family sedans or about sports cars, at this point the MRE 18 would have figured out that

the user's interest falls into a certain category. Based on the user's interest category, the system then retrieves the advertisements that are relevant to that category. Thus, the purpose of the MRE 18 is to figure out what the user requested, to place the user's request in a category of a classification system (i.e., the active index SIC 19) and, based on that classification, to retrieve relevant advertisements. At this point there may be several hundred related advertisements retrieved from the category containing the user's request. This creates a problem in that there will be many more advertisements related to the user's request than can possibly be shown to the user on the first page of information. At this point, the prime space manager 20 comes into play. The purpose of the prime space manager 20 is to prioritize the advertisements that are to be shown to the user.

For example, if there is only room for five advertisements and the MRE 18 returns a list of 500 possible advertisements to show, the system must determine which five of the ads to show. The operation of the prime space manager 20 is based on a set of contracts that were accepted when the advertisers placed their advertisements with the system. The system supports the following five basic types of contracts:

1. Ratio-based: The ad will be shown in prime space no less than a certain percentage of the times a relevant query is made by a searcher.
2. Dependent: Ad A is to be shown with a predetermined probability when ad B is shown in prime space.
3. Exclusive or Competitive: Ad A is never to be shown in prime space when ad B is shown.

4. Story: An ad consisting of several sub-ads will be shown in prime space in a particular sequence to a consumer.

5. Frequency-based: The ad will be shown in prime space a certain number of times in a given time period (e.g., 100 times per month).

Additionally, the contracts may also specify the type of the layout used to mix ads together with the data. This includes the layout types, and so forth. The various contracts used by the system will now be discussed in more detail.

The ratio-based contracts are used where an advertiser wants the advertisement to be displayed a certain percentage of the time in prime space. In other words, when a query is received that is related to the business of the advertiser, the advertiser is guaranteed that its advertisement will be on the first screen of information displayed to the user a certain percentage of the time (e.g., 15% of the time). This percentage is called the coverage guarantee. The ratio-based contract provides a coverage guarantee to the advertiser that its advertisement will be displayed X% of the time for every relevant query.

This guarantee is probabilistic in nature, so that over the long run of the system, the probability will converge to the percentage that the advertiser has requested.

The second kind of contract that the system supports is referred to as a dependent contract. The dependent contract guarantees that one advertiser's ad can be shown in conjunction with another advertiser's ad on the same screen. For example, a travel destination advertisement (e.g., "Come to the Bahamas") can be shown on the first screen of information together with a credit card

advertisement (e.g., "Pay for your ticket through VISA"). With this type of contract, the system can support contracts that satisfy the needs of two or more advertisers at the same time.

5 The third type of contact supported by the system is called an exclusive or competitive contract. This type of contract can be used to guarantee that an advertiser's ad is not shown at the same time or on the same page as a competitor's advertisement. For example, a beverage
10 manufacturer (e.g., Coca-Cola) can prevent a competing beverage manufacturer's (e.g., Pepsi's) ads to be shown at the same time as its own ads and vice versa.

 The fourth kind of contract is called a story contract. This type of contract is used to guarantee an
15 advertising package in which a sequence of two or more ads are shown together. The ads can be displayed so that the user sees a first part of the advertisement, then sees a second part of the ad, and then sees a third part of the ad and so on. For example, an advertisement for coffee might
20 consist of a picture on a first page showing somebody drinking a cup of coffee; a picture on a second page showing them putting the coffee cup away; and a picture on a third page identifying the name of the advertiser's product (e.g., "Maxwell House makes a good cup of coffee").

25 Thus, the story contract guarantees that an advertiser's ads will be shown in the form of a story. The story contract is desirable where a group of ads must be shown together in a particular order to make sense.

 The fifth kind of contract is the frequency-based
30 contract. This contract is used where the advertiser wants to guarantee that its advertisement will be shown a certain

number of times over a given time period (e.g., 100 times per week).

Several different combinations of the above-described types of contracts can also be used. For example, an advertiser can enter a story contract which also guarantees that its advertisement will not be shown at the same time as its competitors ads are shown, that its advertisement will be shown at the same time that a related advertisement is shown, and that its advertisement will be shown in prime space a certain percentage of the time.

The prime space manager 20 has the responsibility of determining all the possible advertisements that can be shown in conjunction with the information requested by the user. The prime space manager 20 then decides which ads to show based on the particular contracts that have been signed. If a contract has already been entered with an advertiser's competitor, the prime space manager 20 will make sure that it can in fact satisfy the new contract with the advertiser. Thus, the prime space manager 20 dynamically makes a decision of which advertisements to show given the amount and the number of advertisements that it can possibly show. The end result of the prime space manager 20's computation is that the system knows exactly which advertisements can and will be shown to the user.

These advertisements are then given to the gate 15. The gate 15 at this point gives the ads to the typography manager 14, and the typography manager 14 determines how best to display the information. The page information then comes back to the gate 15 and is forwarded to the client 17. In the above description of the present invention, the advertisement contracts, the active index SIC 19, the layout rules 11, the segmenter 12, the segment

rules 13, the layout manager 10, and the typography manager 14 have been explained.

During the computation of the advertisements and all the other computations that the system of the present invention performs, a logging module 22 of the system performs extensive logging of what the user has asked, what advertisements were shown, how long the advertisements were shown, and which advertisements were shown to which user. The logging module 22 then stores these logs in a SYS logs database 23. Various scanned reports can be produced and defined using the information in the SYS logs database 23.

A local client 24 of the present invention provides a download delay-time advertising feature. Download delay time refers to the delay experienced by most users of computer networks from the moment a page is requested to the time the page is downloaded to their browser. The download delay-time advertising feature utilizes the time that the consumer would have waited for a page to be downloaded to expose the consumer to ads. Once a page (or a subpart of the page) has been downloaded to the consumer's browser and is available to be displayed, the ads are replaced by the page content.

In most cases, when a user asks for a page there is a considerable download delay time. The user must wait for the page after the request for a few seconds or, in certain cases, a few minutes before the page is finally downloaded by the network. The present invention uses this delay time to show advertisements, while taking care that the delay time of the user is not increased by showing the ad. Thus, the only time an advertisement is displayed is the time that the user would have had to wait anyway. While the

user waits for the page, the user can view the advertisements on the local client 24.

The information brokers or content providers shown in Fig. 1 include a home page dispatcher 25, a search engine
5 INFORMIX 26, and a generic HTML 27. For purposes of the present invention, it is assumed that there are three broad classes of publishers that can utilize the advertising features of the present invention. A "publisher" can include virtually anyone that provides content to the
10 network. For example, anyone who is a home page owner is a publisher in the category shown as Generic HTML 27. A second kind of publisher is the search engine publisher 26, which includes phone company yellow page providers, such as NYNEX. And a third kind of publisher is the so-called home
15 page dispatchers, which include traditional magazines and newspapers, such as Business Week.

The present invention distinguishes between these three types of publishers. If the publisher is a home page dispatcher 25, the publisher provides its own content and
20 intersperses the content with certain special tags indicating where ads can and cannot be placed and indicating where the focus of the content is about certain topics. These special tags are then used to trigger
certain actions in the system of the present invention
25 relating to the type of content contained on the page and the kinds of advertisements that can be placed on the page.

The second kind of publisher that the present invention is used with is the search engine publisher 26. Currently, there are many companies on the WWW that permit
30 users to query their database and then return a set of answers from the database to the user. For example, a telephone company may have a site that allows a user to

obtain a set of phone numbers and business names for a particular type of business (i.e., a yellow page directory).

For purposes of the present invention, the search engine publisher 26 is distinguished from the home page dispatcher 25 in the sense that the content returned by the search engine publisher 26 does not contain any special tags or meta comments put in by the publisher to define the layout of the content and the ads. In this case, the layout manager 10 of the present invention computes the optimum layout based upon the rules and layout templates, as described above. The final result, therefore, is that output is taken from the search engine publisher 26, adorned with certain relevant advertisements, and then shown to the users.

The third kind of publisher, the generic HTML publisher 27, includes those publishers that the present invention assumes are people who have designed their own home page, perhaps taking great care to provide a desired appearance. These publishers will probably not have inserted any special tags or meta comments into their home page to guide the advertising system of the present invention. In this case, the system of the present invention can either try to find a space in the home page that it thinks is reasonable without violating the desired appearance of the page, or it can insert the advertisement before or after the page. That is, if the segmenter 12 cannot find a suitable place within the page that the generic HTML publisher 27 has already designed, it will place the advertisement on the proceeding or the following page. This is more likely to occur if the owner of the page has not provided any special tags or meta comments, or

if copyright concerns prevent the insertion of a certain advertisement within the page.

In summary, the present invention partitions the publishers (i.e, the information brokers) into three
5 categories: the home page dispatchers 25, the search engine publishers 26, and the generic HTML home page providers 27.

The distinction between the three is that the home page dispatchers 25 intersperse their content with special tags or meta comments which helps the present invention in
10 determining where to place the advertisements and what kind of advertisements to place, the search engine kind of publishers 26 return unadorned results of queries and the present invention is free to reformat and relay out the set of answers, and the generic HTML home page publishers 27
15 are publishers that have meticulously crafted their home pages, usually without special tags or meta comments, so that no guidance is provided to determine an appropriate space for placing an ad.

The segmenter 12, which is the module of the present
20 invention whose responsibility is to find a place somewhere on a generic HTML page to place an advertisement, may not succeed in finding a space on a generic HTML home page. In this case, it will either place the advertisement on a proceeding or a following page. Most, if not all, of the
25 publishers on the WWW will fall into one or more of these three categories.

In Fig. 1, the acronym LA refers to List of Ads, Q refers to the user's query, CGI is a standard WWW
30 abbreviation for "common gateway interface," HTML is a standard abbreviation for hypertext markup language, and Ly is the list from the layout manager 10 to the typography

manager 14 that indicates a list of pages that have to be displayed.

Referring to Fig. 2, a flow chart is provided that describes the main flow of the process of placing ads on content pages provided by yellow page publishers, in accordance with the present invention. The process begins at START 30 when the client provides a query to the system.

For purposes of illustration, the query depicted is a yellow page query in which a user requests some item of interest, and the publisher is a yellow page publisher (please refer to the previous discussion for a description of the three kinds of publishers).

For example, the user may submit a query form asking for a listing of all car dealers in Morristown, New Jersey. This query form is evaluated by the box 31 entitled YP QUERY FORM. The purpose of this evaluation is to determine the type of the query (e.g., is the user interested in cars, a particular geographic region, etc.).

In using a yellow page publisher there are two broad distinctions for a query. A client may be asking for a certain category of listings, or the client may be asking for a particular vendor. For example, the user could ask for car dealers in Morristown, NJ (i.e., a category of listings), or the user could ask for Morristown BMW located on South Street in Morristown, NJ (i.e., a particular vendor). The system determines which of the two types of queries or searches the user has made, as illustrated by box 32 in Fig. 2. If the query is for a certain category, the process will go to the left hand side of the flow chart of Fig. 2, and if the query is for a certain vendor, the process will go to the right hand side of the flow chart of

Fig. 2. The left hand side of the flow chart will be explained first.

After determining the type of query, the category search engine 33 next determines which category best fits the user's request. The user may have asked for "car," but the category in the yellow page provider's index may in fact say "automobile." Or, the user may have asked for "spectacles," and the category in the yellow page provider may be called "optician." The matching of these variations of terms is performed by the category search engine 33.

Once it has been determined which category the user's request falls into, the advertisement selection process comes into play with the ad selector 34. The ad selector 34 determines what advertisements are best suited to be mixed in with what the user has requested. The content from the category search engine 33 and the ad(s) from the ad selector 34 are then given to a mixer 35. The mixer 35 functions to mix the content coming from the search engine with the ad(s) selected by the ad selector 34. The result is the creation of a page that is of interest to the user.

If the user's category was about cars, at this point the ad selector 34 would have presumably found advertisements related to cars and displayed these ads at step 36 to the user. The user at this point can select one or more of the listed or advertised car dealers. This again results in more advertisements being selected by the ad selector 34, as well as the accessing of a vendor search engine 37 provided by the yellow page publisher, and a new page being shown to the user.

Referring to the right hand side of Fig. 2, if the user requests a particular vendor, the system will first do a search time estimation 38 to figure how long it will take

to find the piece of information requested by the user. At
the same time, the ad selector 34 will perform a search for
the advertisements that are to be shown to the user. As
soon as the relevant ads are found by the ad selector 34,
5 the ads are mixed with the search time estimation
information and both are displayed to the user. At this
point 40 in the process, the user is asked for a
confirmation of whether to continue with the requested
search given the estimated search time. Upon receiving
10 confirmation from the user, the ad selector 34 will again
be accessed, as will be the vendor search engine 41. The
first set of ads returned by the ad selector will continue
to be displayed to the user for viewing until the search
results from the vendor search engine 41 are available.

15 When the vendor search engine 41 completes the search,
it returns the search results to the mixer 35. Similarly,
the ad selector 34 returns a second set of ads to the mixer
35. The mixer 35 mixes the ads with the content provided
by the vendor search engine 41 and forwards the combined
20 page to the page display 42.

The result of this process is that a yellow page
inquiry posed by the user is categorized either as a
category search or a vendor search. In both cases, the

25 type of the query is matched with a certain category of
ads, the search results and the ads are mixed, and a page
is constructed and shown to the user. Upon further input
from the user, the system selects and mixes a second set of
ads with the content page of the publisher, and another
page containing ads is constructed and shown to the user.

30 The "focus" arrows 43 shown in Fig. 2 indicate that a
certain focus is associated with each category. The query
may have been directed to a category of listings or a

particular vendor. In both cases there is a "focus" associated with the content of the query (e.g., automobiles, physicians, lawyers, etc.). In addition, there may be a focus associated with the geographic location of the user to permit advertisers to target users in particular geographic regions. The focus process plays a major part in the present invention. No advertisements are shown unless it can be determined that the advertisements are in some way focused or related to the content of what the user requested.

Referring to Figs. 3A-3D, the page views displayed to the user during the process of Fig. 2 will be described. Figs. 3A-3D show four form views labeled successively Form# YP_1, Form# YP_2, Form# YP_3, and Form# YP_4. These form views correspond to the references made to the same form numbers in the flow chart of Fig. 2.

The user is first presented with a user's screen that looks like Form# YP_1 and the user is prompted to type in a query. For example, the user can type in a query for car dealers in Morristown, NJ. This query will result in a search type as previously described, either as a category or a vendor. In fact, the user can pick one or the other by selecting the appropriate box on Form# YP_1. At this point, the system knows whether the search type is category or vendor and, because the user has typed in car dealers, the system knows the focus.

At this point, if the search type is vendor, the system will show a screen that looks like Form# YP_2. This screen informs the user that it will take approximately 30 seconds for the system to respond to the user and asks the user whether this is okay. If the user wants to proceed,

the user clicks on the box labeled "Search it now" and the system begins the search.

If the search type is category, the system displays a screen that looks like Form# YP_3. This Form# YP_3 shows a list of categories under the heading of car dealers. The user is then asked to make a selection of one of these categories. Simultaneously, the system goes into the ad selection mode and selects and places ads on the Form# YP_3 that is being displayed to the user. In this case the advertisement shown on Form# YP_3 is the Bellcore logo.

The user then selects one of the subcategories being displayed and clicks on the box "Search it now."

The system next displays a screen that looks like Form# YP_4. The ads placed on this Form# YP_4 have been selected by the ad selector 34, and the listings on the form have been searched and retrieved by the vendor DB entries search engine 41.

In summary, the general process for a yellow page search request in the system is that first the user states a query, the result of the query is a set of possible categories that the user wants to see. The system then makes a choice of one category from this list of categories and shows the relevant providers in that category to the user. When the user states a query and the system displays to the user a list of categories, certain advertisements are shown to the user as well. This is the top half of the flow chart of Fig. 2 where the category search engine 33 or search time estimator 38 and the ad selector 34 provide the two components to the mixer 35 and the mixer 35 comes up with a page containing both an advertisement and a list of categories or the search time estimation of how long it will take.

If a user requested a category, the user will be shown advertisements and a list of categories that relate to the user's query. If the user requested a particular vendor, the user will be shown advertisements related to the focus of the user's query and an estimation of how long it will take for the system to process the query. In both cases, the user has been shown advertisements and some part of the information and then asked to make a further choice. Once the user makes the further choice, the system again tries to find what the user requested plus what advertisements to show as a part of this process. Thus, whenever the user asks the system to do something, the system will do what the user asks, but at the same time will find advertisements that it can mix in with what the user requested.

Although the flow chart of Fig. 2 illustrates the present invention only in conjunction with a yellow page publisher, similar processes are used with other types of publishers. The process of the present invention for placing ads with home page providers is described below in conjunction with the flow chart of Fig. 7.

Referring now to Fig. 4, the layout manager 10 of the present invention will be described in further detail. At step 50, the layout manager 10 is given the collection of data to be displayed. The layout manager 10 has a collection of rules that it evaluates so that it can return the best layout for each page. When the layout manager 10 receives data it reads the parameters of the data for each rule (step 51) and expands on the expand meta template 53.

There are two main components of the layout manager 10, as previously discussed -- a layout data base and a layout rule base. The layout rules describe certain

parameters that can be used to determine the best layout to
use under a given set of situations. The layout data base
is a collection of templates that can be thought of as
forms that need to be filled out with certain kinds of
5 - data.

When the layout manager 10 is given a piece of data,
it uses the rules to determine which templates to use.
Specifically, the layout manager 10 evaluates each and
every rule in its rule base to figure out which rule is the
10 best for a given piece of data. The rules are evaluated at
step 54 by calculating a meta and a target function cost.
The meta template is the template that describes what the
layout is about. For example, the template may provide
that there are only five places for advertisements in a
15 particular layout.

There is a cost associated with each kind of layout.
The basic problem that the layout manager 10 is faced with
is that there are a certain number of advertisements to show and
there is a certain amount of user-requested data that has
20 to be shown. This creates a spectrum of objective versus
subjective content that has to be evaluated. That is, it
must be determined how much subjective content to provide
and how much objective content to provide. If the
publisher is a yellow page publisher, the layout manager 10
25 must decide how many telephone listings to display and how
much of the page to fill with advertisements. The system
could, for example, display only three telephone listings
and fill the rest of the page with advertisements, or the
system could display twenty telephone listings and only one
30 ad. Each template in the layout database has a set of meta
attributes associated therewith indicating the number of
advertisements that can be displayed, the amount of content

to be displayed, and the cost function associated with the template.

In summary, the layout manager 10 picks up a rule in the data base, calculates the cost function for each
5 template and stores the template with the minimum cost in a safe place (step 55). At the end of this computation the layout manager 10 succeeds in finding the minimum cost layout within its data base that it can use for the current data account. It then picks the minimum cost layout as the
10 best layout and returns the best layout (step 52) to the typography manager 14 for display to the user.

Referring to Fig. 5, a flow chart of a parser 60 of the present invention will be described. The purpose of the parser 60 is to interpret the meta comments and meta
15 tags inserted on the publishers' pages. As described above, these are generally pages constructed by magazine or newspaper publishers that have provided special tags in their content for the present system to interpret in a special way. These pages are referred to in Fig. 5 as
20 enhanced HTML or E_HTML 61.

The parser 60 first searches for special tags (step 62) in the input E_HTML. There are two kinds of special tags: tags for focus and tags for advertisement space.

Tags for focus are tags that tell the system what the focus
25 of the page is or what the context or content of the page is about. For example, a focus tag might indicate that the page is about sports cars, cooking, or travel destinations in the Caribbean. The advertisement space tags, on the other hand, indicate a space on the page that is so many
30 pixels by so many pixels and that can be used for advertisements. The parser 60 searches for these special tags and identifies the tag name as either a focus or an

advertisement tag (step 63). If the parser 60 cannot find any special tags, it simply returns from where it was called.

5 If the parser 60 finds a focus kind of tag, it extracts and stores the focus words (step 64). If the parser 60 finds an advertisement kind of tag, it extracts and stores the size value (step 65). In both cases, the process returns the focus and ads information to another module in the system (step 67) after reaching the end of
10 the page (EOF) 66.

Referring to Fig. 6, a process flow of the mixer and ad selector will be described. The purpose of the mixer 35 (as previously described in reference to Fig. 2) is to take publishers' content and advertisements and combine them
15 together so that the content and the advertisements are mixed on the same page.

In Fig. 6, the mixer 35 is shown receiving two inputs from the publishers: data 50 (which is the content) and E-HTML 61 (which contains the special tags). The layout
20 manager 10 and parser 60 both form a part of the mixer 35. The data 50 is input to the layout manager 10, and the E_HTML 61 is input to the E_HTML parser 60, as previously discussed. Both of these sub-modules then determine where
the advertisements can be placed on the publisher's page.

25 The advertisement list is then input from the ad selector 34. The ad selector 34 receives a focus input 43, retrieves relevant ads (step 70), and creates the advertisement list using the prime space manager 20 (step 71). These advertisements are then placed in the parser 60
30 and the layout manager 10 (step 72), as described above. The mixer 35 then logs all the essential billing and other user information (step 73) for keeping track of the

system's placement of an advertiser's ad. At this point, a
refresh tag is inserted (step 74) and the system outputs an
HTML page (step 75).

The HTML page output from the mixer 35 has no special
tags and can be displayed by any client that can understand
HTML. The E_HTML 61 is peculiar to the present invention.

The HTML 75 is the standard HTML used on computer
networks. The E_HTML 61 is peculiar in the sense that it
has the special tags for advertisements and focus, as
discussed above.

The refresh tag inserted at step 74 is a special tag
inserted into every page that is to be displayed. The
purpose of this tag is to refresh the page if the user does
not take any action for a predetermined period of time.

The system will refresh and show a new advertisement in the
ad space, while leaving the rest of the page the same.
Thus, while the user is viewing the page, the system will
not show the same advertisements for more than a
predetermined period of time so as to take advantage of the
user's limited attention span. This permits the system to
effectively run as many ads as possible.

For example, while a user is reading an online
magazine page, the advertisement may change every so often
as the user is viewing the same page. If the user is
making changes to the page or is going back and forth on
the page, then the advertisements will not be changed by
the refresh tag function. The timer for the refresh tag
function can be specified by the advertiser or determined
by the system.

As described above, there are five basic types of
contracts supported by this system. One type of contract,
the ratio-based contract, requires that a certain

advertisement has to be shown on the first screen of information or the prime space screen of information for a certain percentage of the time. In order to satisfy this contract, the system has to be able to refresh the prime space screen. Otherwise, once an advertisement was placed on a user's screen it would remain on the screen the entire time that the user is viewing the screen. The insert refresh tag function 74 ensures that the user will be shown new advertisements and not be subjected to the same advertisement for any great length of time.

Referring to Fig. 7, a flow chart of the present invention as used with a home page provider will be described. The flow chart of Fig. 7 is somewhat similar to the flow chart of Fig. 2. However, while the flow chart of Fig. 2 shows the process used with yellow page publishers, the flow chart of Fig. 7 shows the process used with home page provider publishers. As described above, home page providers are publishers that have provided special tags and special focus and advertisement space tags in the meta content supplied to the present system.

To start (step 80), the user enters a query. For example, the user may enter restaurants or cars as a query.

The query has a focus, as described above. The system determines what the focus is and, as described above, the system provides the user with a list of categories that relate to the query. For example, if the user requests restaurants, the user might be shown a list of restaurant types, such as Chinese, American, French, Italian, and so forth. The query entered by the user is evaluated by a query form manager (step 81) to determine the focus of the query.

The screen ("Form# 1") illustrated in Fig. 8A is the graphical interface used to enter the query. Form# 1 shows the space (immediately following the phrase "SEARCH FOR") where the user types in the query (in this case "car").

5 The result of the query is a list of relevant subcategories for cars, as shown in Fig. 8B ("Form# 2").

Referring again to Fig. 7, the category search engine 33 generates a number of subcategories of the query (e.g., cars), and the ad selector 34 (Fig. 7) finds certain

10 advertisements that are related to cars. The mixer 35 mixes these two screens to create the display shown as Form# 2 in Fig. 8B (step 82).

At this point (step 82) the user is asked to make a further selection. That is, the user is asked to pick one

15 or more of the categories displayed in Form# 2 so that the system will provide more specific information. The user may pick, for example, "new and used car dealers retail" by highlighting that item and clicking on the submit button. Upon clicking the submit button, the content provider

20 search engine 83 is accessed. At the same time the ad selector 34 is directed to find appropriate advertisements for the category "new and used car dealers retail."

The content provider search engine 83 and the ad selector 34 then return their respective outputs to the

25 mixer 35. The mixer 35 then mixes the outputs and produces a list of matching content providers plus the advertisements (step 84). The combined display at this point is shown as Form# 3 in Fig. 8C. Form# 3 shows a list of different car dealers in the center of the display, and

30 five advertisements along both sides of the display.

The user is then further asked to make a selection of a particular dealer (e.g., a Mazda 626 dealer or a Ford

Mustang dealer) that the user wants to visit. The user makes the choice, and the choice of the user is given to the content provider dispatcher 85 and the ad selector 34.

As the content provider dispatcher 85 is retrieving the requested home page, the ad selector 34 again finds an appropriate advertisement based on the user's request.

For example, the user can pick a Mazda 626 dealer by clicking on Mazda 626. As the page for a Mazda 626 dealer is being retrieved, the ad selector 34 finds advertisements that can be inserted on this page. There may or may not be advertising space on the home page for the Mazda 626 dealer.

If there are no spaces on the home page for the system to place ads, the system will either insert the ads before or after the page, or the system may decide not to place any advertisements at all. In any case, for any home pages requested, the ad selector 34 will attempt to find suitable advertisements and give these ads to the mixer 35, and the mixer 35 at that point will display the content provider's home page with or without ads.

In summary, Figs. 8A-8C show three screen displays referred to in Fig. 7 as Form# 1, Form# 2 and Form# 3, respectively. Form# 1 is the first screen that the user sees and in which the user is expected to type in a general query in the typeable space. This results in a second screen (Form# 2) where the user is asked to make a choice from a list of categories. After choosing from the list of categories in the second screen, a third screen (Form# 3) is shown to the user from which the user makes a choice from a list of subcategories. The present invention uses these three screens also to show ads related to the user's requests. Form# 1 shows two ads for the owner of the

service provider (e.g., Bellcore in this case). Form# 2 shows advertisements related to cars because the user typed in cars. Form# 3 shows more car advertisements relating to particular makes of cars because the user asked for retail car dealers.

Referring to Figs. 9A and 9B, a description of the billing system of the present invention will be provided. Figs. 9A and 9B show two submodules of the system called the info bank and ad auditing. The purpose of these two modules is to construct log reports of various usages of the present invention.

Fig. 9A shows the process flow for the info bank submodule 90. The purpose of the info bank submodule 90 is to log certain kinds of information during a particular client's session. The logged information is then used to provide advertisers with a description of what advertisements have been shown, what the user has bought after seeing an ad, and so forth. Thus, the info bank submodule provides an information bank to show advertisers who has been accessing the system and how successful the ads have been.

The screen for the relevant module contains a link called the info bank which the user can select. When the user clicks on the info bank link, the user is shown the start 91 of the info bank submodule. At this point the user is asked for authentication information (step 92), such as the user's name and personal identification number (PIN). The system then goes into the authentication mode 93 and verifies whether the user ID and PIN number typed in are valid. If the authentication is successful, the system then searches the transaction log and filter (step 94) and pulls up the appropriate log requested by the user. The

user can ask to see all transactions performed by the client or all of the things of a certain type that the user has bought. A report of the results (step 95) is automatically produced for the client and displayed.

5 The submodule for ad auditing 96 has a similar purpose to the info bank submodule. With the ad auditor, an advertiser that has placed an advertisement with the system is able to go into the system at any point in time and find out who has seen his ad, what was the content of the query
10 when the advertisement was displayed, how many times has the advertisement been displayed, and so forth. The authentication and other process steps are similar to the corresponding steps of the info bank submodule 90.

15 More specifically, in order to use the ad auditing submodule 96 the advertiser is shown a screen with a link for ad auditing. The user clicks on the ad auditing link and is led into the start 97 of the ad auditing module. The advertiser is then asked to type in authentication information (step 98), which consists of the user name
20 followed by a PIN number. If the PIN and the user name check out and are authenticated by the authentication step 99, the user is then asked to tell the system which advertisement or which subset of advertisements they want to verify or audit (step 100). Once the user specifies the
25 advertisement by name or by browsing all the advertisements in the system and clicking on the ones that the user wants to get a report on, the system generates a report (step 101), which is a log of who has seen the ad, what machine they were using, what time and for how long the
30 advertisement was displayed, and other desired information.

All the log information accessed by the info bank and ad auditing submodules is logged in a single database 23

(Fig. 1). Whether the user is looking for auditing information, transactions, things that the user has bought, or user demographics, the information is all in the same database. However, when a user uses the info bank
5 submodule 90, the user is interested in only certain things, as described above. But when the user uses the ad auditing submodule 96, the user is interested in other things. Therefore, the reports that are produced by the info bank and ad auditing submodules contain different
10 data. The database is filtered and only the items of interest for a certain report are retained so that the rest of the database need not be used.

Referring to Fig. 10, the flow of an ad placement process 110 according to the present invention will be
15 described. The purpose of ad placement is to allow advertisers to enter their advertisements into the system.

For entering an ad, the system provides a screen that is shown to the user asking whether the user wants to enter an ad. If the user indicates yes by clicking on that
20 particular choice, the system enters the start 111 of the ad placement mode. At this point the system asks the user for the focus (step 112). The advertiser may say, for example, that he is in the car business, the car washing
business, or that he is a physician, a lawyer or whatever
25 other category name that he wants to give. The user is also asked for an advertisement name at step 112. This is just a name for future reference.

The purpose of the focus in step 112, as discussed above, is to prevent an advertisement from being shown that
30 is not relevant to the query at hand. The system of the present invention always shows advertisements that are relevant to what the user has asked for. Therefore, it is

of paramount importance that the system know the context of the ad. Thus, when the advertiser places an ad, the system establishes the focus.

5 The present invention describes the world in terms of SIC codes, which is a categorization of approximately 64,000 categories. For example, the system knows that there is a distinction between opera glasses and operating tables. It knows a distinction between lawyers and physicians, as well as the distinction between different
10 kinds of physicians. For example, when the user types in bicycle as the focus, the present invention accesses the category search engine 33 and returns a list of subcategories related to the bicycle business (e.g., bicycle repair, bicycle sales, bicycling as a sport, and so
15 forth). The user is prompted (step 114) to select one or more of the subcategories from the list that are relevant to the user's ad. The user can also create his or her own categories at this point. The result is that the advertisement is tied to all of the categories selected or
20 created by the user.

The user may also be asked to provide certain demographic or profile information. For instance, the user can require that his advertisement be shown only to people in age group 30 to 40 or only to people living in
25 Morristown, NJ or any other geographic location. The last item that the user is asked to specify is the contract. The various contracts available to the advertiser are explained above. When the user is finished entering all of this information, the system updates the ad info database
30 (step 115).

The system then either accepts the advertisement (step 116), or the system rejects it. The system might reject

the advertisement because the contract that the user wants cannot be satisfied by the system. If the system decides at this point that the user's contract cannot be satisfied, the system displays to the user the next best possible contract that it can satisfy, and the user is then free to either choose that contract or to suggest a different one for consideration by the system. Upon confirmation of the ad placement, the advertisement becomes a part of the database.

Referring to Fig. 11, the flow of a page placement process 120 according to the present invention will be described. The purpose of the page placement process is similar to the ad placement process described above with reference to Fig. 10. The page placement process 120 is distinct from the case where the page resides outside of the present system, either as a part of a search engine or because somebody else owns the home page. The page placement process is used for publishers interested in creating a home page for use in the present system. The page placement process permits publishers to provide a page with meta tags and meta contacts.

The page placement process has a starting screen 121 in which one of the menu choices is, "Would you like to place a page?" If the user clicks on that menu choice, a screen is shown to the user asking the user to type in a page name, a page URL (i.e., Universal Resource Locator), and a focus. The URL is a physical address in terms of the WWW of where the page is actually residing. As discussed above, the focus is a list of key words that the user believes captures the content of the page.

Upon entry of the focus, page name, and page URL, the system retrieves a list of subcategories from the category

search engine 33 that are related to the focus at hand and displays those subcategories to the user (step 124). The user picks one or more of these subcategories as being relevant to the content of the page. The system is then
5 able to integrate the page into the page database of the present invention (step 125) and return a confirmation of the new home page being linked to the relevant categories (step 126).

Referring to Fig. 12, an architecture of the
10 underlying transaction system of the present invention will be described. One of the basic features of the present invention is that the advertisements displayed to users provide a gateway for the client to take further action with the advertiser. That is, the ads are not just there
15 for the user to see, they are interactive.

For example, if the system shows the user a picture of a car or a cookie and the user wants to buy or receive further information about the advertised product, the user can make a transaction with the system. The transaction
20 can be very simple, such as where the user clicks on the advertisement and is displayed all information about that particular product. Or the transaction could be something more complicated, such as where the user actually enters a dialogue with the system and buys the product.

25 The flow chart shown in Fig. 12 describes a typical scenario in which the client has seen a certain advertisement and clicks on that advertisement to actually buy something. In particular, whatever this person is going to buy is going to involve multiple servers. The
30 client is running on Host 1, and the transaction that the client is requesting will involve multiple servers, perhaps geographically distributed. Host 4 is a bank server

running at a bank somewhere on the computer network. Host 3 is a transaction server for logging and serving as the transaction focal point. Host 2 is a server that accepts the start of the transaction from the client housed at Host 1.

A typical transaction where the client 130 sees an advertisement and clicks on the advertisement will now be described by reference to the transaction system architecture diagram shown in Fig. 12 and the screen displays shown in Figs. 13A and 13B. The action of clicking on the advertisement is captured by the gate 131 in Host 2, and the appropriate information is then displayed to the client by the transaction client 132 in Host 2. For example, the transaction client 132 might prompt the client to enter the type of transaction desired.

The transaction client in Host 2 is a client for a transaction server 133 in Host 3.

Referring to Figs. 13A and 13B, in the typical transaction the user first clicks on an item from a list of features offered by the particular publisher (Fig. 13A), in this case, the item "Internet DMV." A second screen display (Fig. 13B) then appears in which an advertisement is provided in the middle of the page that says "Division of Motor Vehicles, Morristown, New Jersey, Automatic Renewal of Driving License." The user may then click on this particular advertisement. At this point a screen is displayed to the user asking the user what type of transaction the user wants to make. The next screen generated by the transaction client at Host 2 would say, for example, "Would you like to renew your license?" or "Would you like to get a new car registration?" If the user says, "I would like to renew my driver's license," the

request is captured by the transaction client 132 at Host 2 and given to the transaction server 133 in Host 3. The transaction server 133 in Host 3 then receives this request for renewal of driving license, extracts the appropriate forms to be filled out from its local database, and provides the forms to the transaction client 132 in Host 2.

The transaction client 132 at Host 2 then receives these forms and fills them out with the client's information with help from the client 130 at Host 1. Thus, the client's name, address, car ID and other necessary information are filled in for the client automatically.

The client 130 is asked to type in a PIN number to start the second part of the transaction. If the client 130 types in the PIN number, the second part of the transaction starts. The second part of the transaction is that the form that has now been filled in and constitutes the request for a renewal of driver's license is sent from the client 130 in Host 1 to the transaction client 132 in Host 2, and from the transaction client 132 in Host 2 to the transaction server 133 in Host 3. At this point, there are two things that have to happen at Host 3. First, the particular request for renewal of the driving license has to be processed by making a record that this person has renewed the license and by sending confirmation back to the client. Second, the transaction server 133 must be paid for the renewal of the driving license. In other words, there has to be a monetary fee that has to be exchanged between the client and the DMV.

In exchanging the monetary fee, the transaction server 133 will also act as a transaction client 134 to the bank server 135 in Host 4. The transaction client 134 sends a message to the bank server 135 in Host 4 indicating that it

has the client's PIN number and the authorization to deduct a certain fee from the client's account. The bank server 135 authenticates that information, logs the information in its own records, and sends back an electronic approval for the debit.

The transaction client 134 at that point indicates receipt of the payment to the transaction server 133 in Host 3. The transaction server 133 in Host 3 logs receipt of the money, receipt of the request, and sends a confirmation message through the transaction client 132 and gate 131 back to the client 130 at Host 1. The confirmation might say, for example, "We have received your request; we have deducted \$17 from your account; your license is in the mail; here is your new driving license number; thank you very much."

Referring to Fig. 14, the transaction process of the present invention will now be described by reference to a flow chart. The user starts the process (step 140) by clicking on an appropriate advertisement to ask for a transaction. The system then links the user (step 141) to the transaction server (step 142). The transaction server then returns the required form to the client (step 143), and the client fills in the form with the user's profile information. The form is then resubmitted to the transaction server (step 144). The transaction server accesses the bank server (step 145) to determine whether the user has the required money and whether the money can be withdrawn.

Upon authentication by the bank server, the transaction server receives confirmation from the bank server (step 146) that the user has given a valid account number, and that the bank server has in fact authenticated

the user and has authenticated the deduction of the appropriate sums from the bank. The transaction server then makes the deduction and puts the deducted amount into the bank server's log (step 147). These logs are
5 maintained so that at any time users can find out how much money they have spent or who has bought what things. The transaction server also maintains a transaction server log of the transaction time and any other particulars of the transaction desired (step 148). Once the logging is
10 complete, the transaction server then generates a confirmation of the transaction which is then displayed to the client (step 149).

Referring to Fig. 15, the prime space manager 20 will now be described in further detail. As explained above, an
15 important feature of the present invention is that it offers advertisers certain contractual guarantees as to how and when their ads will be shown in prime space (the first screen of information displayed to the user). These contractual guarantees include frequency-based contracts,
20 ratio or percentage-based contracts, exclusive or competitive contracts, dependent contracts, and story contracts. These contracts can be used alone or combined in a variety of ways, as previously explained.

The prime space manager 20 has a prime space algorithm
25 150 that accepts as inputs the list of relevant ads that are related to the query at hand, the size of the prime space, and the list of all ad contracts. The prime space manager 20 then looks at all the ad contracts and determines which advertisements to show to the user. The
30 prime space manager 20 then updates the contracts (step 151) to indicate which ads were selected (this information is used in considering which ads to select the next time).

The prime space manager 20 then provides a list of prime space ads to be shown to the appropriate module. That is, the prime space manager 20 accepts a possibly very long list of advertisements and, based on the contracts that it has, filters the ads to a smaller list containing exactly the number of advertisements that it can show based on the size of prime space.

The following examples of prime space algorithms are provided to further illustrate the function and operation of the prime space manager 20 of the present invention. It should be noted that many variations of these algorithms are possible, and that the specific examples provided are for purposes of illustration only.

Notations:

Let A_1, \dots, A_n denote the advertisements in the system of the present invention.

Let P_1, \dots, P_n denote the contract values associated with the corresponding advertisements (P_i denotes the contract value associated with advertisement A_i , as a function of the price paid by the advertiser.)

Let S denote the size of the prime space in the system.

Prime Space:

The prime space on a user's terminal means the first screen of data which constitutes the answer to the user's query (plus the relevant advertisements). In this context, 5
advertisers will be assured of "fair" treatment in the sense that their advertisements will be presented to the targeted users in a manner that is commensurate with the fees they were charged. Fairness means equal access as well as controlled access--ensuring that advertisers who 10
pay more are more accessible than those who pay less.

Contracts:

Contracts are specifications that indicate the frequency of display of a given advertisement in response 15
to relevant user queries.

The example prime space algorithms provided below may be used for the following four types of advertisement contracts:

- 20 1. Frequency-based: These contracts specify the number of times a given advertisement is to be shown in response to the relevant user queries.
- 25 2. Ratio-based: These contracts specify how often a given advertisement is to be shown in the prime space in response to the relevant user queries.
- 30 3. Dependent: An ad A_i can have a dependent contract of $X\%$ with another ad A_j . The dependent contract will ensure that ad A_i is shown together with A_j at least $X\%$ of the times ad A_j will be shown.

4. Exclusive: An ad A_i can have an exclusive contract with another ad A_j . The exclusive contract will ensure that ad A_i will not be displayed along with ad A_j .

5 **Algorithm for Frequency-based Contracts:**

WHILE (screen space is not filled)

begin

1. Select an advertisement, say A_i , with the selection biased by the contract value P_i

10 2. Include A_i in the screen space

end

Algorithm for Ratio-based Contracts:

WHILE (prime space is not filled)

15 begin

1. Select an advertisement, say A_i , with the selection biased by the contract value P_i

2. Include A_i in the prime space

end

20

Algorithm for Dependent Contracts:

WHILE (Prime space is not filled)

begin

1. Select an advertisement, say A_i , with the selection biased by the contract value P_i

25

2. Assign A_i to A

3. WHILE (some advertisement is selected AND prime space is not filled)

- Select one of the dependents of A, say A_j , with the selection biased by the contract value P_j

30

WO 97/21183

PCT/US96/19509

- Assign A_j to A
- Include A in the prime space

end

5 **Algorithm for Exclusive Contracts:**

WHILE (prime space is not filled)

begin

1. Select an advertisement, say A_i , with the selection biased by the contract value P_i
- 10 2. Remove all advertisements A_j 's that have an exclusive contract on A_i
3. Remove all advertisements A_k 's that A_i has an exclusive contract on
4. Include A_i in the prime space

15 end

Algorithm for Overall Prime Space Management:

WHILE (prime space is not filled)

begin

- 20 1. Select an advertisement, say A_i , with the selection biased by the contract value P_i
2. Assign A_i to A
3. WHILE (some advertisement is selected AND prime space is not filled)
 - 25 - Remove all advertisements A_j 's that have an exclusive contract on A
 - Remove all advertisements A_k 's that A has an exclusive contract with
 - Select one of the dependents of A, say A_m , with the selection biased by the contract value P_m
- 30

WO 97/21183

PCT/US96/19509

- Assign A_m to A
- Include A in the prime space

end

5 **Algorithm for Checking Consistency:**

Notations:

Let P_A denote the prime space contract value of Node A .

Let $D(A,B,X)$ denote the dependency contract value from advertisement A to advertisement B (A should appear together with advertisement B , $X\%$ of the times advertisement B appears).

Let $E(A,B)$ denote the exclusive contract from advertisement A to advertisement B (note that no value is associated with exclusive contracts).

15 If a new advertisement A is to be introduced into the system with a prime space contract value of P_A , a set of dependency contracts $D(A,1,x), \dots, D(A,k,y)$ and a set of exclusive contracts $E(A,1), \dots, E(A,q)$, then for checking the consistency of all the contracts, the following basic checks have to be carried out (more specific checks may be required for specific algorithms and their implementations):

1. Check whether Inequality (I) can be satisfied
2. For each of the dependency contracts $D(A,i,x)$ check whether Inequality (II) can be satisfied
- 25 3. For each of the exclusive contracts $E(A,i)$ check whether conditions (III) can be satisfied

30 The consistency checking algorithm verifies if the new advertisement can be accommodated with the corresponding contract. If the new advertisement can be accommodated, the contract is approved and the system is updated to

reflect the new addition, else, the consistency checking algorithm informs the advertiser of the maximal contract that can be accommodated for the new advertisement.

Install Equation Editor and double-click here to view equation.

Install Equation Editor and double-click here to view equation.

For checking the Dependency Contracts and the
5 Exclusive Contracts a graph is drawn by representing all the advertisements as nodes in a directed graph. A dependency contract "Advertisement C appears 50% of the times as advertisement D appears (together)" is represented as an edge from the node " C " to the node " D ". A weight of
10 0.5 is assigned to this edge. For each exclusive contract $E(C,D)$, there exists an edge in the graph. Thus, a graph is formed where the nodes denote the advertisements, and the edges denote the dependency and exclusive contracts.

If a new dependency contract $D(A,B,X)$ is to be
15 introduced, then the following check has to be made for consistency:

Install Equation Editor and double-click here to view equation.

If a new exclusive contract $E(A,B)$ is to be introduced, the following conditions (III) should hold (for
20 any value of x):

Install Equation Editor and double-click here to view equation.

where $X_{(A,B)}$ denotes the weight of the dependency edge from A to B , i.e., $D(A,B,X_{(A,B)})$.

Fairness:

The fairness guarantees of an algorithm for the system ensures that the advertisements are displayed in a way that is proportional to the fee charged for that advertisement. The fairness of the frequency-based contract algorithm, ratio-based contract algorithm, and the dependent contract algorithm follows from the probability distribution of the contract values and the random number generator that selects an advertisement with the correct random bias. The probability that the random number generator will generate a number in a particular range is proportional to the length of the range. The larger the range, the more probable that a number corresponding to that range will be selected. Each advertisement is assigned a range of values depending on the fee charged for that advertisement. The more the fee charged, the greater the range assigned to that advertisement. This ensures that the larger the fee, the more frequently (more number of times) that advertisement will be displayed.

Exclusive contracts are constraints imposed on the display of advertisements. Fairness in meeting the exclusive contracts means that the contracts are satisfied (no two advertisements that have an exclusive contract are displayed together). The prime space management algorithm ensures that the exclusive contracts are met, by removing all the advertisements that have a relative contract with an advertisement A, as soon as advertisement A is selected for display in the prime space.

In order to ensure fairness, the contracts have to be consistent. When contracts are negotiated with the advertiser, certain consistency checks (using the

consistency checking algorithm) are made in order to ensure that the fairness guarantees can in fact be met.

Story Advertisements:

5 Story advertisement is a kind of advertisement in which the advertisement changes dynamically as time progresses. The story advertisement is composed of a sequence of static conventional advertisements, which together tell a story. A story advertisement is displayed
10 using the notion of refresh-insert tags. A refresh tag that is associated with a page denotes the time period after which the story advertisements are refreshed (the next advertisement in the sequence is displayed). All other static advertisements are unaffected by the refresh
15 timer. The refresh tag can be set by the client or the application. Thus, the story advertisement is a logical sequence of a number of advertisements, that will be displayed one at a time at each refresh cycle.

 It will be appreciated that the present invention is
20 not limited to the exact construction or process steps that have been described above and illustrated in the accompanying drawings, and that various modifications and changes can be made without departing from the scope
25 thereof. It is intended that the scope of the invention only be limited by the appended claims.

WHAT IS CLAIMED IS:

1. A system for advertising on a computer network, comprising:

a server containing a plurality of advertisements;

5 means for electronically connecting the server to a computer network; and

means for selecting and retrieving an advertisement from said server in response to a query entered on said network.

10

2. The system for advertising as set forth in claim 1, wherein said selecting means comprises means for ensuring that a selected advertisement is relevant to the query.

15

3. The system for advertising as set forth in claim 1, further comprising a mixer means for combining a retrieved advertisement with a content page returned by the computer network in response to the query.

20

4. The system for advertising as set forth in claim 1, further comprising a prime space manager means for controlling when the advertisements contained on said

25 server are selected by said selecting means, said control being based upon respective advertiser contracts associated with each of said advertisements.

5. The system for advertising as set forth in claim 1, further comprising a local client means for displaying a retrieved advertisement only during a download delay time period in which a user is waiting for a content page to be downloaded in response to the query.

30

6. The system for advertising as set forth in claim 1, further comprising means for calculating a search time estimation for the query, and means for combining a page
5 containing said search time estimation and a retrieved advertisement.

7. The system for advertising as set forth in claim 1, further comprising means for determining a geographic
10 location of a user that entered the query, and means for selecting an advertisement based upon the user's geographic location.

8. A method for advertising on a computer network, comprising the steps of:

15 placing a plurality of advertisements on a server;
placing the server in electronic communication with a computer network;

selecting an advertisement on said server in response to a query entered on said network; and


20 retrieving the selected advertisement from said server for display to a user.

9. The method for advertising as set forth in claim 24, further comprising the step of combining the retrieved advertisement with a content page returned by the computer
25 network in response to the query.

10. The method for advertising as set forth in claim 24, further comprising the step of controlling when the advertisements on the server are selected based upon
30 respective advertiser contracts associated with each of the advertisements.


11. The method for advertising as set forth in claim
24, further comprising the step of displaying the retrieved
advertisement only during a download delay time period in
which a user is waiting for a content page to be downloaded
5 in response to the query.

FIG. 1



2/17

FIG. 2



WO 97/21183

PCT/US96/19509

3/17



FIG. 3A

FORM# YP_1




FIG. 3B

FORM# YP_2



FIG. 3C

FORM# YP_3




FIG. 3D

FORM# YP_4

FIG. 4

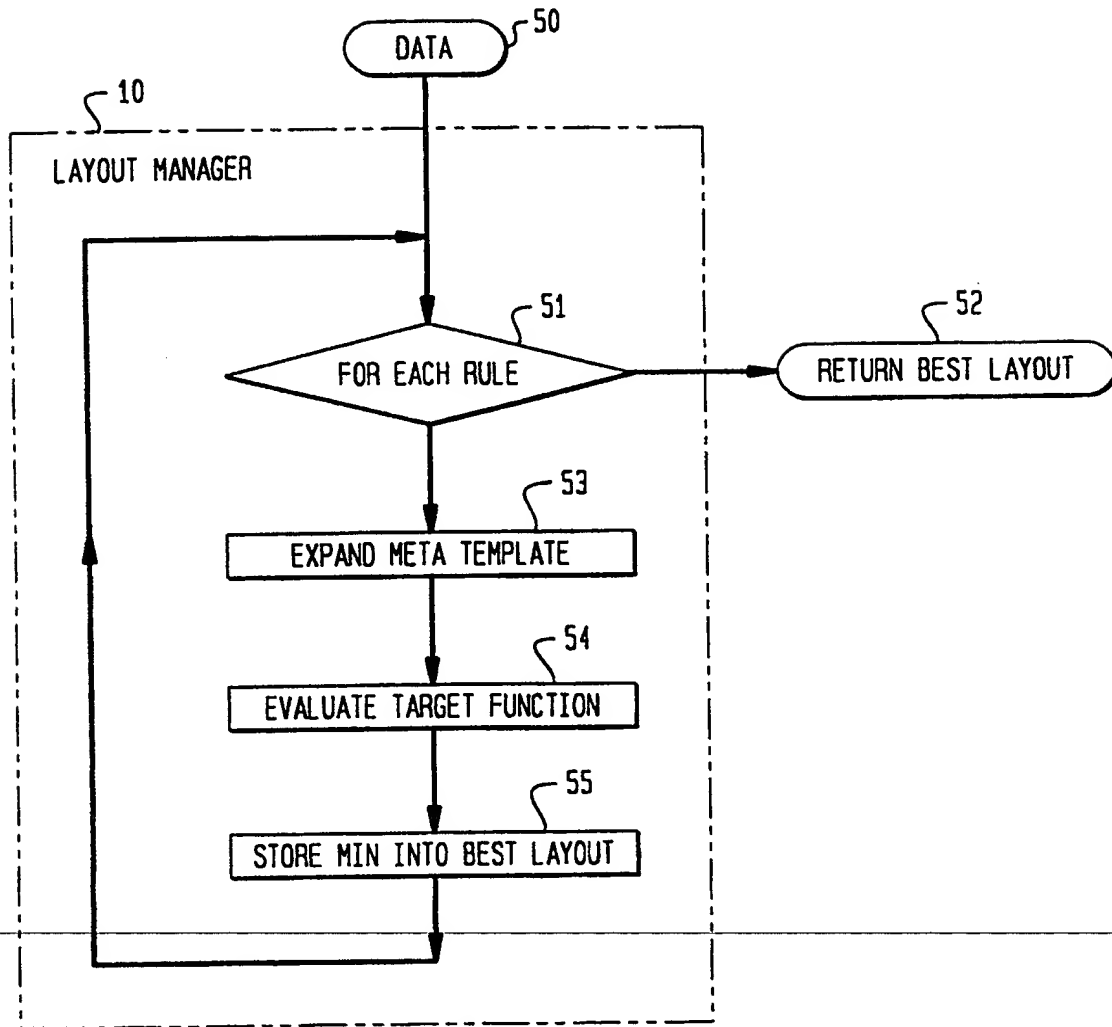


FIG. 5


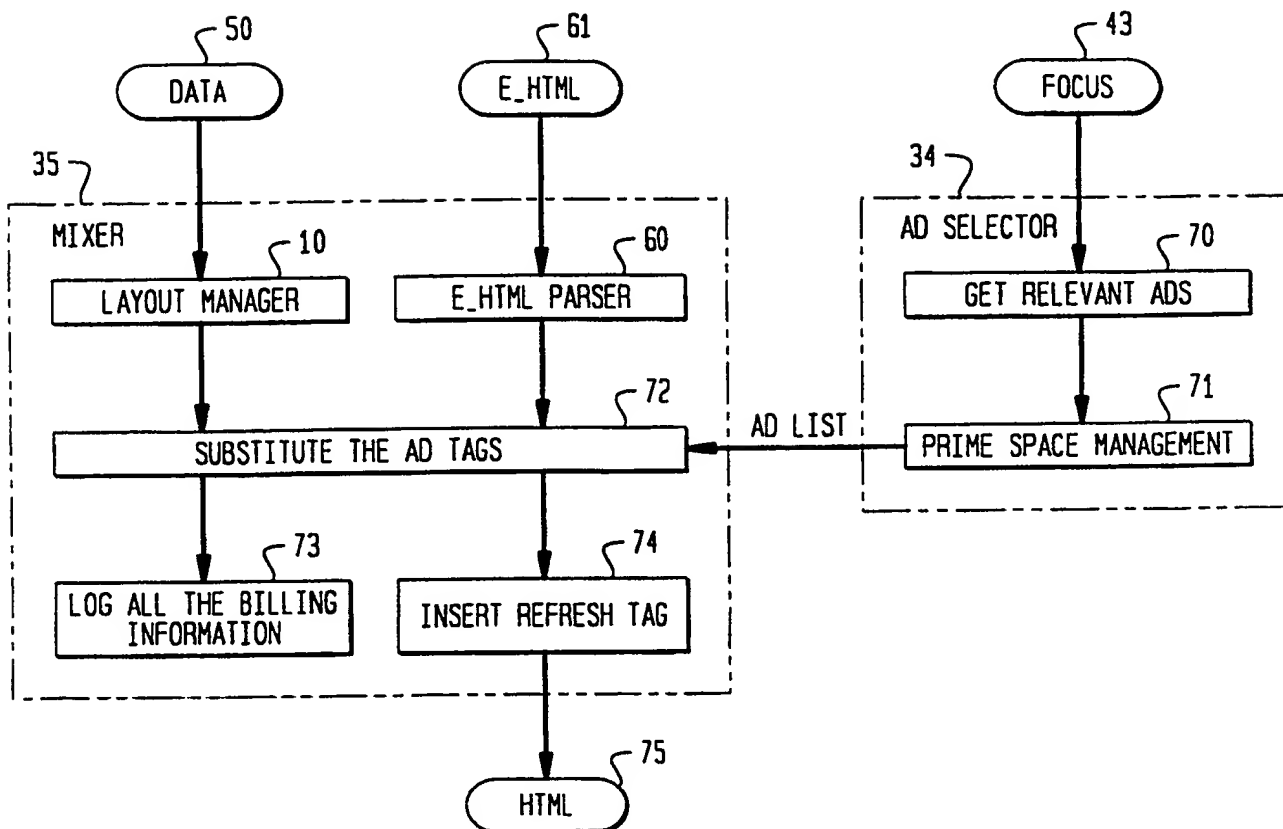
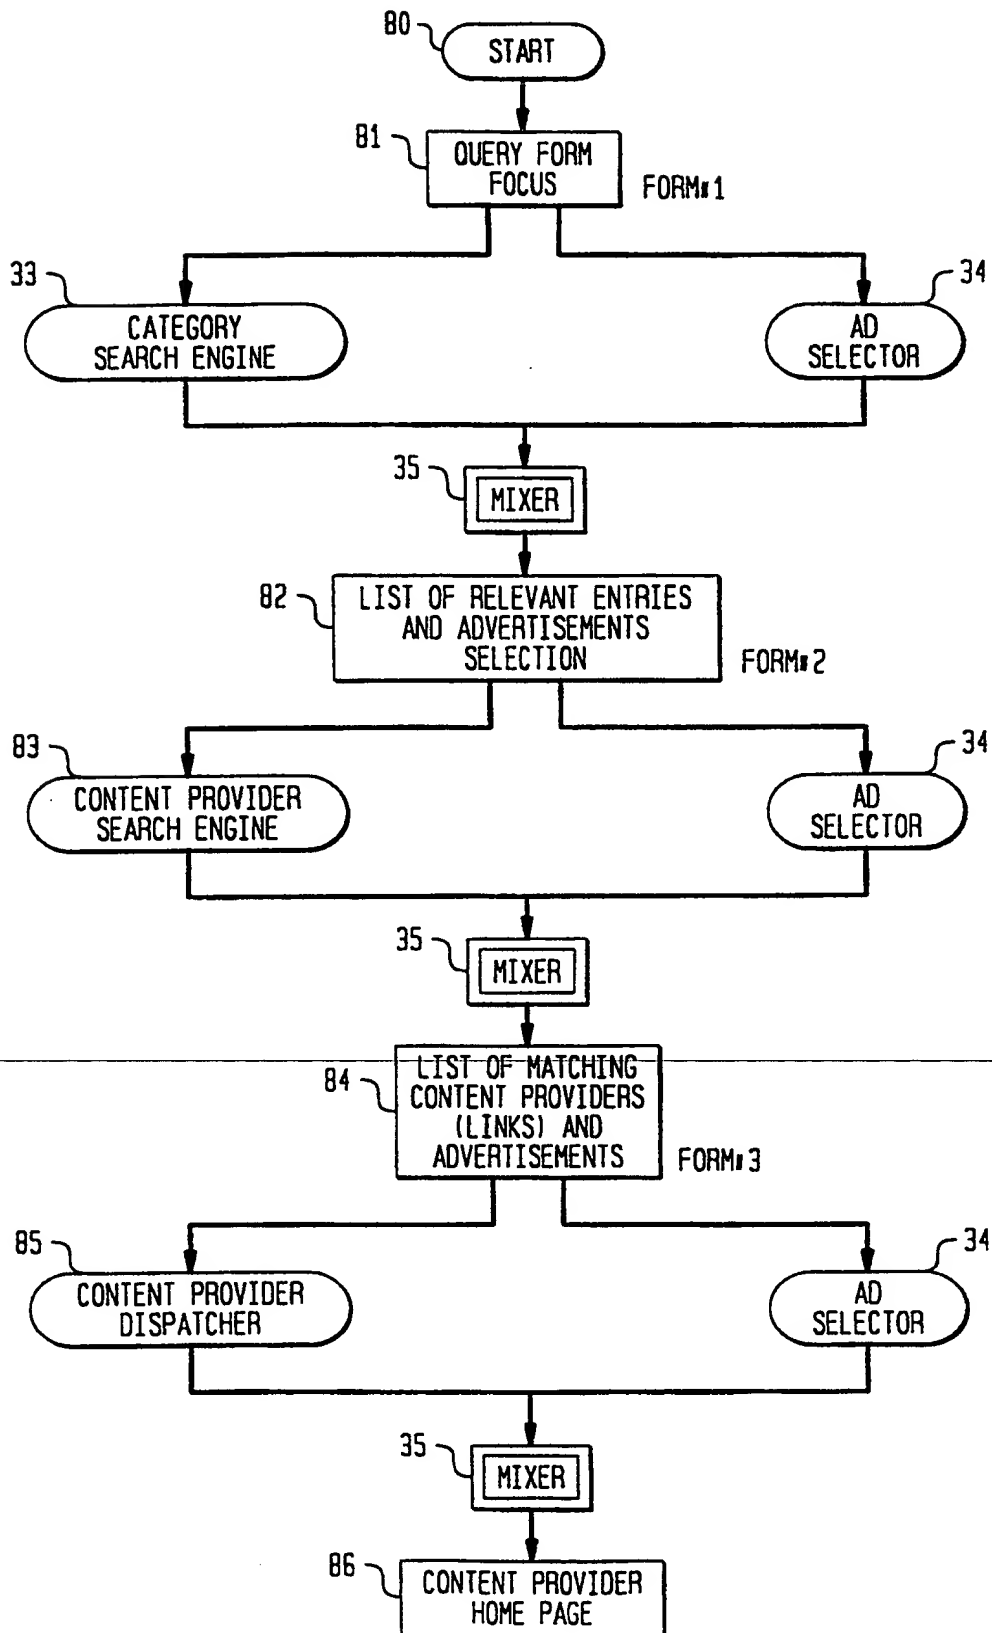


FIG. 6



7/17

FIG. 7




WO 97/21183

PCT/US96/19509

8/17

FIG. 8A




WO 97/21183

PCT/US96/19509

9/17


FIG. 8B



PCT/US96/19509

10/17

FIG. 8C



SUBSTITUTE SHEET (RULE 26)

FIG. 9A




FIG. 9B




FIG. 10

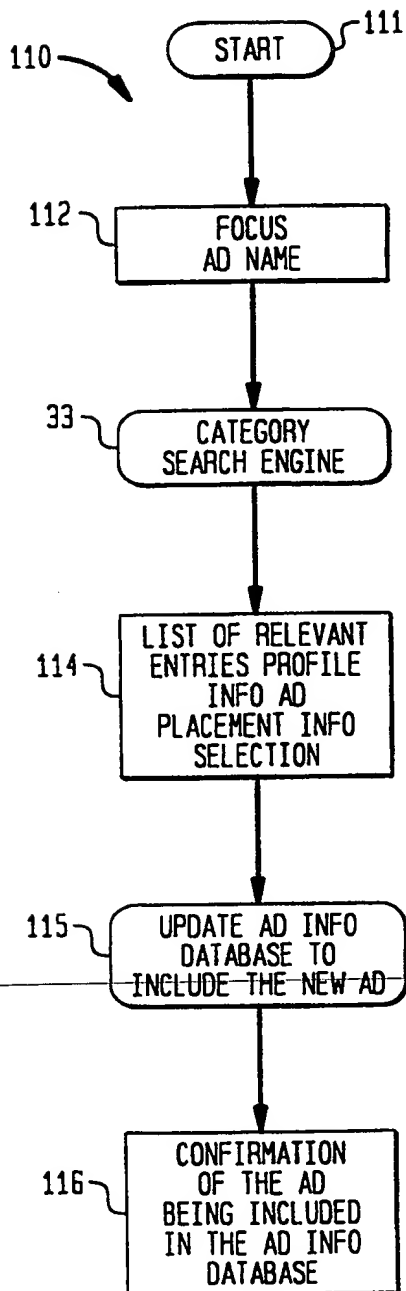


FIG. 11





FIG. 12




WO 97/21183

PCT/US96/19509

14/17

FIG. 13A

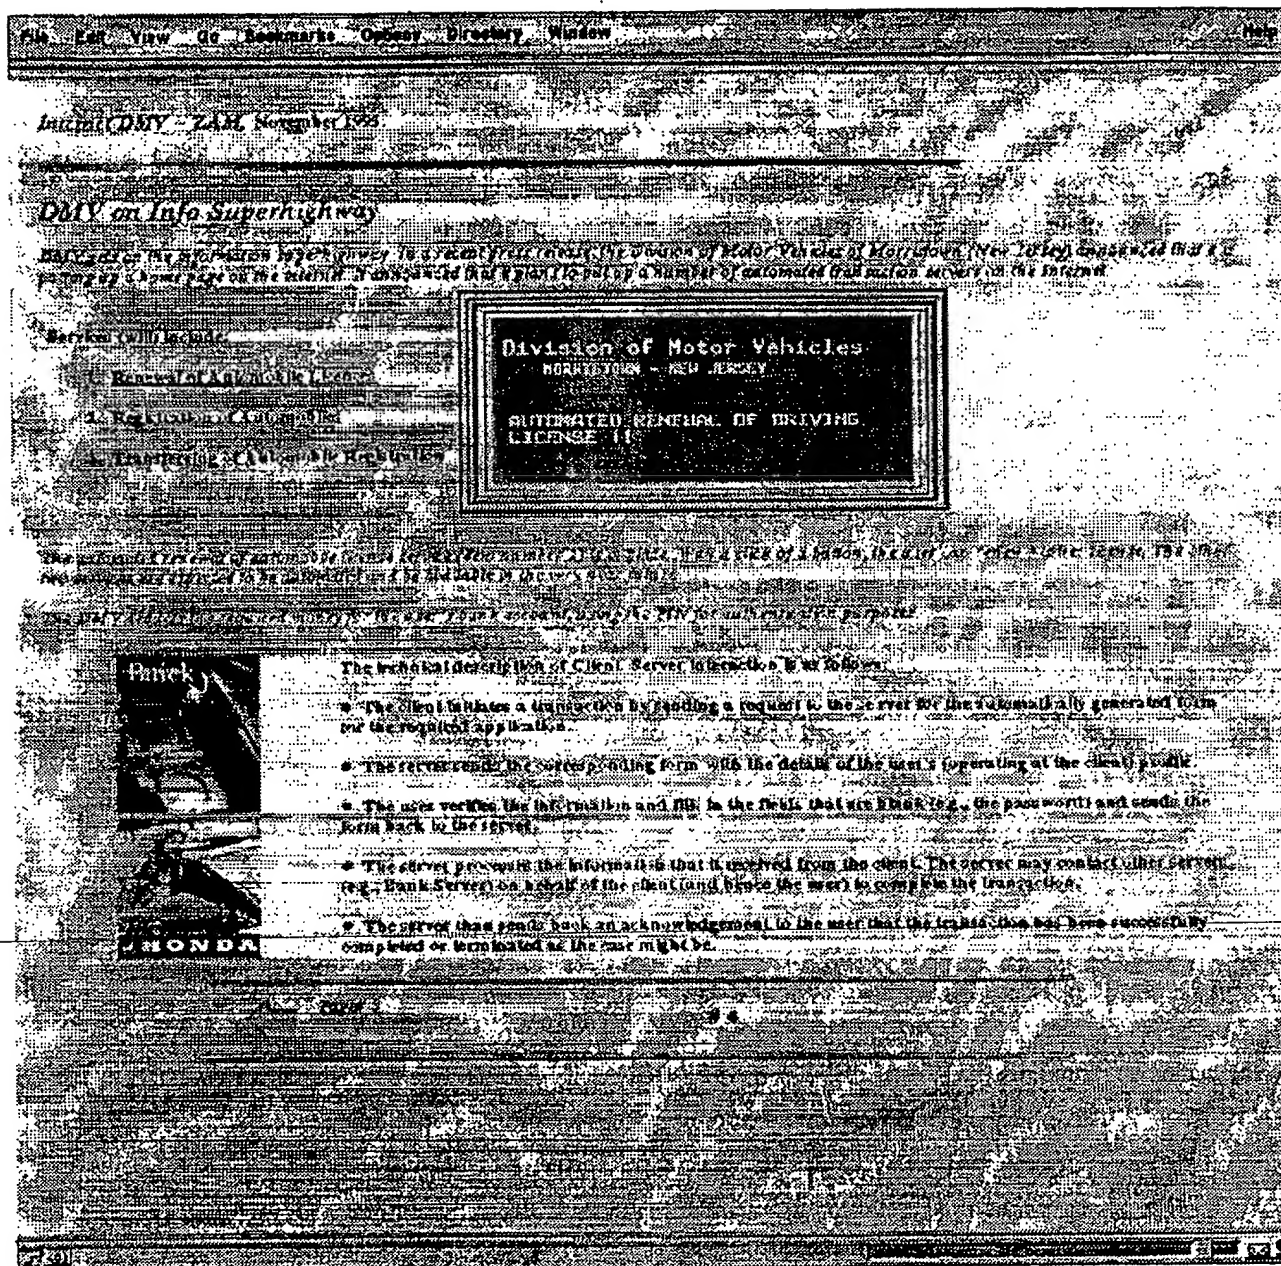


WO 97/21183

PCT/US96/19509

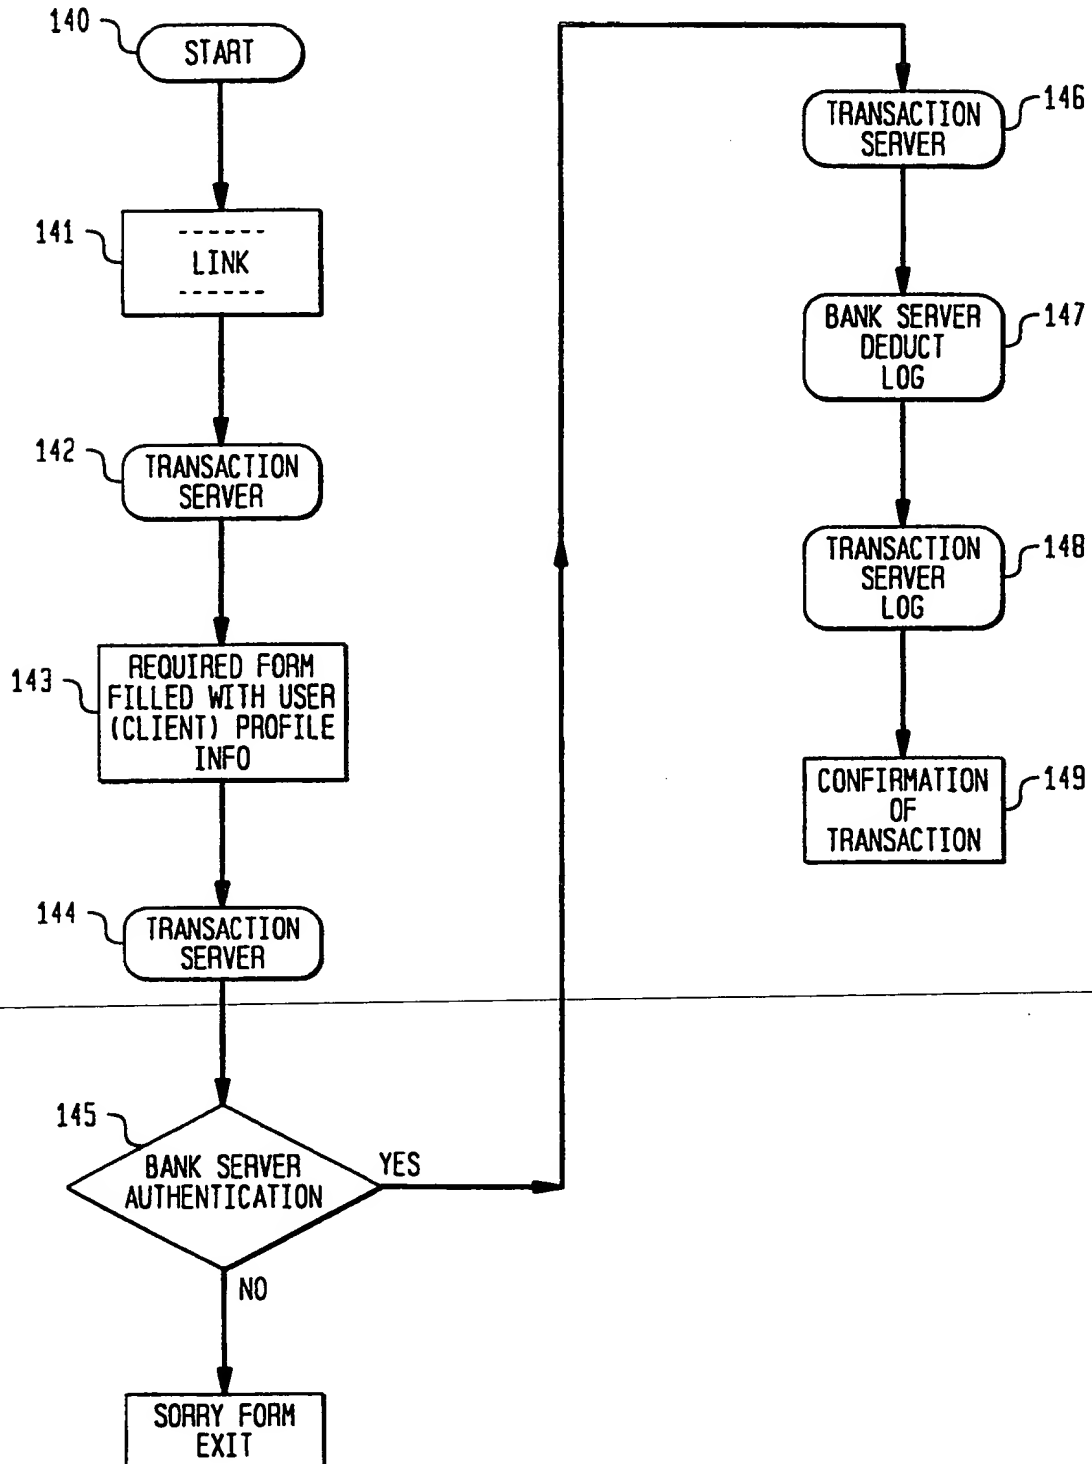
15/17

FIG. 13B



16/17

FIG. 14



17/17

FIG. 15

